Detecting and Locating Rogue Access Points
Eric Junker
CprE 537, Iowa State University
Spring Semester 2003
Supervisor: Dr. Steve Russell
April 28, 2003

This report has been erroneously attributed to me (Steve Russell) because Eric and I agreed that I could use it as a report example in subsequent classes. In internet searches, my name was mistakenly attributed. The report was written by Eric and future attributions should be given to Eric.

Detecting and Locating Rogue Access Points

Eric Junker
CprE 537
Spring 2003
Dr. Steve Russell

April 28, 2003

Abstract

The problem of rogue access points is becoming a problem that can lead to serious security vulnerabilities in a network. This paper will discuss several methods of detecting rouge access points such as: TCP Fingerprinting, SNMP Scanning, Packet Sniffing, Active Probing and RF Monitoring. Once a rogue access point has been discovered it is important to locate the rogue access point so that the appropriate measures can be taken. Location methods such as: Time of Arrival, Angle of Arrival and Received Signal Strength are discussed.

Contents

List of Figures

Chapter 1

Introduction

From a security point of view, wireless LANs represent a new method for accessing the enterprise network. This new method of network access introduces some new security problems that need to be solved. Many papers have been written on securing wireless networks against unauthorized attacks, not much attention has been paid to detecting and tracking these breaches once they occur.

With proper care, it is possible to build a wireless network that is at least as secure as an equivalent wired network. However, it is still desirable to track attempts at unauthorized access and, perhaps more importantly, track potential breaches by internal, trusted users.

The purpose of this paper is to discuss techniques for detection and location of rogue access points and then to discuss the advantages and disadvantages of each approach.

1.1 Rogue Access Points

A rogue access point is simply an access point that is not accounted for. It's not officially part of the network. It transmits and eventually becomes accepted as part of the network. The major danger is that on most networks, once you are part of the network there is no further protection scheme in place to limit access to network resources. So once you're in, you're in, and you have access to all the data that's on the network. In some cases companies have mistakenly put their wireless LANs inside the firewall so in effect once you're in the wireless LAN you're in the corporate wired LAN and have access to everything including routers. The problem of rogue access points may not seem to be a big problem but the Gartner Group estimates that 20% of enterprises have rogue wireless LANs attached to their networks.

Rogue access points deployed by end users pose a great security risk. End users are not security experts, and may not be aware of the risks posed by wireless LANs. Any user can purchase an access point, and connect it to the corporate network without authorization. Little does this user know that the user has unwittingly opened a gaping hole in the local network, such that any drive-by attacker could simply hop on the local network and do anything they wish. Some departments may install their own wireless networks without proper authorization. Most access points do not have their security features enabled and many have not had any changed made to the default settings. To make matters worse, groups of wireless enthusiasts and war-chalkers are documenting and publicizing the locations of open access points, greatly increasing the likelihood that this security breach will be exploited.

The Pentagon was recently found to have insecure APs, even though they were managed by the IT department. Chris O'Ferrell, chief technology officer at NetSec Inc. in Herndon, Va., which provides intrusion-detection services to numerous federal agencies and commercial customers, detected the non-secure wireless LAN at the Defense Information Systems Agency (DISA) on May 10.

While parked across the street from DISA's headquarters, O'Ferrell was able to view the Service Set Identifier (SSID) numbers of access points and numerous IP addresses. Using a standard 802.11b wireless LAN card attached to his laptop computer and AP detection software from NetStumbler.com, he was able to scan the network in less than half an hour.

Chapter 2

Detecting Rogue Access Points

As stated previously rouge APs can pose a serious security threat to wireless networks. An ideal solution would be able to detect such behavior and send a response such as sending relevant information such as geographical location and the time when the element was detected to the network administrator. It is also possible in some situations to disable the rogue AP so that it can't be used.

If you can obtain the MAC address of the rogue AP you can query the Ethernet switch to determine the switch port each rogue access point is connected to. In theory, once you discover the switch port, the Ethernet cable connected to that port could be traced back through a series of patch panels to the office or cubicle where the AP resides.

Practically this method doesn't work well. First the method is time consuming. You may have to query many switches to find the one that is being used by the rogue access point. Secondly, once you find the switch that the rogue access point is connected to, you have to trace the cable from the wiring closet to the office or cubicle that the rogue access point is located in. This process may be impossible if there aren't accurate up to date wiring diagrams and maps. Another problem is that APs have both radio and LAN MAC addresses, which differ from each other, and I can detect only the radio address. I need some way to match up that radio address to the LAN MAC address I can see on the switches.

There are two main categories of methods to detect rogue access points. The first category of methods are those that use an existing wired LAN to scan for the rogue APs. Such methods include TCP fingerprinting, SNMP scanning, and sniffing. The second main category of detecting rogue APs is to examine the radio link. Examples of this type of detection are active probing and RF monitoring.

2.1 Wired Based Detection Methods

2.1.1 TCP Fingerprinting

TCP fingerprinting is the process of examining the subtle differences in how a target responds to various specially crafted packets in hopes of determining the OS of the target system. Basically, you just look for things that differ among operating systems and write a probe for the difference. If you combine enough of these, you can narrow down the OS very tightly. You can use TCP fingerprinting to identify rogue APs by scanning your entire network and looking for hosts that have match the OS profile of an 802.11 AP. See Figure A.1

The advantages of this method is that you can simply start up a tool such as Nmap that performs TCP fingerprinting and have it scan your entire network while you go out to lunch. Once you return from lunch you can analyze the results to see if it found any rogue APs.

The disadvantages are that if you have a large network it could take a long time to scan your entire network. It is also not the best method because the act of determining the OS is intrusive, noisy and is likely to look like suspicious activity to IDSs. The other disadvantage is that TCP fingerprinting is not 100% accurate. Nmap uses the information to make a guess at the OS but there are bound to be times when it has false positives or false negatives.

2.1.2 SNMP Scanning

SNMP scanning can also be used to detect rogue APs. SNMP scanning is similar to TCP fingerprinting but instead of using the differences in the TCP/IP network stack it uses information obtained by using the SNMP protocol. If a host as UDP port 161 open it is likely that it is running the SNMP service.

As seen in Figure A.2, you can use the command snmpwalk to extract information about the target host. From the example scan we can see that the target is manufactured by Cisco which could lead us to believe that it is a Cisco wireless AP.

The advantages of this method are similar to that of TCP fingerprinting in that you can start a scan and let it run while you work on other things. Once the scan is complete you can examine the results. One advantage that SNMP scanning has over TCP fingerprinting is that SNMP scanning is likely to be more accurate than TCP fingerprinting.

The main disadvantage of SNMP scanning is that not all APs support SNMP and it may be turned off which would make it impossible to use SNMP to get information on the device.

2.1.3 Packet Sniffing

The third method of detecting rogue APs from the wired network is to configure a device to run in promiscuous mode and analyze the packets and examine the Ethernet headers to check that the MAC addresses are authorized MAC addresses. Another method is to compare MAC and ARP entries, and look for ports with multiple connections or to compare MAC and ARP entries and look for popular WLAN vendor MAC addresses.

An example of a tool that to monitor MAC addresses is Arpwatch. Arpwatch is a tool that monitors Ethernet activity and keeps a database of Ethernet/IP address pairings. It also reports certain changes via email. As mentioned earlier you can use the MAC address to trace back to the AP or else tell the Ethernet switch to block traffic to the offending MAC address.

The advantage of this method is that it is a continuous process and it is constantly monitoring the network for unauthorized MAC addresses. Once a rogue AP is connected to the network it can be detected once it transmits any data.

The disadvantages are that sniffing the traffic may be considered intrusive since it requires that all the traffic is examined. There is also the problem of scalability. If you have a high speed network it may be hard to process high volumes of traffic and perform the analysis. And finally, it may be possible to evade the detection system by changing the MAC address of the AP to a MAC address of an authorized client.

2.2 Wireless Based Detection Methods

2.2.1 Active Probing

The active probing method uses probe request frames on each channel where it is able to detect wireless activity. When an AP comes within range of the client and receives a probe request frame it will typically respond with a probe response frame containing the network ESSID.

Tools like NetStumbler allow network administrators to wander around looking for unauthorized access points, but it is difficult to devote time to wandering the building looking for new access points.

The advantages of using the active probing method are that it is the easiest method to implement and is presently the only network discovery method used on Windows systems.

The disadvantages include that the person trying to detect the rogue APs must walk around the building with a laptop or a handheld device which is time consuming and expensive.. Another problem is that is may also pick up other access points in the area, which may be a concern if you are sharing a building or a floor with another organization. Their access points may cover part of your floor space. The periodic walk through is the only way to check for unauthorized access points. Active probing will not be able to discover rogue APs that are configured not to advertise their ESSID using a cloaked ESSID configuration or if the AP is configured to not respond to probe requests. They are also easy to elude, since a rogue AP can easily be unplugged when the scan takes place.

2.2.2 RF Monitoring

RF monitoring is a completely passive method of wireless LAN discovery known as radio frequency monitoring (RFMON). A client with a wireless card that is configured in RFMON mode will be able to capture all RF signals on the channels to which it is configured to listen.

This method can detect rogue APs by specifically monitoring raw 802.11 frames to detect if there are any telltale frames broadcast by a rogue access point.

There are three classes of 802.11 frames. If we are trying to detect rouge access points we are primarily interested in class 1 and 2 frames. Class 1 frames are the only frames allowed in state 1, the unauthenticated state and are largely management frames used for authentication, beacons and probe requests. Class 2 frames are allowed in both states 1 and 2 and are used for association and re-association. From rogue access points we would expect to see a large number of beacon frames (Class 1). When a WLAN client encounters an area with multiple Access Points, the client can use the Beacon to determine which AP has the strongest signal to associate with. A Beacon is the announcement that an AP is alive, and is broadcast every 10 seconds.

A detection system could listen for beacon frames and each time a beacon is sniffed, the agent could compare the source MAC address of the frame with a list of registered access points. If not found in the list of registered access points the agent would consider the AP to be a rouge one and send a message such as an email notification to the network administrator. A quick check in the network switches and routers may help to determine where the AP is located. Utilizing vendors signal strength utilities to determine the general area an AP is installed

An advantage of RFMON is that it doesn't have the limitation of only getting information from probe request and probe response frames. An advantage over wired based scanning methods is that there should be far less wireless traffic to analyze than there would be if you were to analyze Ethernet traffic on the wired network.

One of the disadvantages of RF monitoring is that for RFMON to work the client must be in range of the rogue AP so like active probing you must walk around your building with a laptop or handheld device. Another disadvantage of RFMON is that currently free tools that use RFMON scanning are only available on Linux and BSD-based operating systems, with limited support on BSD hosts. Some commercial products are available for Windows, but at a significant cost. Also while in RFMON mode, wireless clients are unable to transmit any frames, their cards are only able to receive, and therefore capture traffic. This limits the client from probing the discovered AP for SNMP MIB information.

2.3 Detection Experiments

2.3.1 TCP Fingerprinting

To perform TCP fingerprinting I used Nmap[1]. I simply had Nmap scan a block of IP addresses and perform TCP fingerprinting. The results of the test can be seen in Figure A.3. Nmap was not able to give an accurate guess of the OS but by looking at the results it can be determined that the target is an AP. The hostname of the target is `airport8` from which we can conclude this is an Apple AirPort wireless AP. Also by looking at the ports that are open such as `DHCP` and `SNMP` we can conclude that the target is a wireless AP.

2.3.2 SNMP Scanning

To perform the SNMP scanning I used the snmpwalk tool on Linux. A simple script could be written to call the snmpwalk command with range of IP addresses and store the results to a file. For my test I simply used snmpwalk to query the AP that I found in the previous section. The results are shown in Figure A.4. One piece of information that sticks out is the line that shows `system.sysDescr.0 = Base Station V3.81 Compatible`. This tells us that the target is a wireless base station. Not all APs have SNMP enabled so this test will not work in every situation.

[1] http://www.insecure.org

2.3.3 Packet Sniffing

For the packet sniffing test I used a program called AirSnare[2]. AirSnare is a program that sniffs packets on a wired or wireless network and monitors the MAC addresses. You can specify a list of known MAC addresses and if an unknown address is found it will alert you. Another way to use AirSnare would be to look for MAC addresses that have the same vendor code as a wireless AP. A screenshot of AirSnare can be seen in Figure A.5.

2.3.4 Active Probing

To perform the active probing test I used Netstumbler[3]. Netstumbler works by sending out probe request frames and then listens for probe reply frames from wireless APs. A screenshot of Netstumbler can be seen in Figure A.6. I was able to find several APs and I was able to see what channel they were using and if WEP was enabled.

2.3.5 RF Monitoring

Kismet[4] was used for the RF Monitoring test. Kismet works by passively listening to 802.11b signals and can detect wireless APs. Since Kismet does not transmit any data it is impossible to detect if someone is using Kismet. The main disadvantage of Kismet is that it only runs under Linux and it doesn't have an easy to use graphical interface. A screenshot of Kismet

[2] http://home.attbi.com/ digitalmatrix/airsnare/

[3] http://www.netstumbler.com

[4] http://www.kismetwireless.com

can be seen in Figure A.7.

2.4 Detection Summary

There are several different methods to detect rogue access points which are shown in Table A.1, each with its own advantages and disadvantages. The RFMON method of detecting rogue APs seems to be the best method overall.

Current research is being done to create a system that uses distributed agents placed throughout a building. The agents scan their respective cells for rogue APs by sniffing for beacon frames emitted by the APs. Each time a beacon is sniffed, the agent compares the source MAC address of the frame with the list of registered access points. If not found in the list of registered users, the agent considers the AP to be a rouge one and sends a message such as an email notification to the network administrator. Along with the alert, the system could provide critical information such as MAC address, SSID, channel, security settings, etc., to enable the IT department to identify, locate and disable the potential rogue immediately. Future versions would also include a tool that could be used to home in on an access point based on real-time signal strength meters. There is also research being conducted which is attempting to discover rogue access points though the use of physical layer RF characteristics.

As wireless networks become more common it will be crucial for every network administrator to know how to scan his network for rogue APs.

Chapter 3

Locating Rogue Access Points

Nothing could be worse than a user installing an unsecured AP in a conference room that is broadcasting into the parking lot, allowing an attacker to gain direct access to the network from outside the building. In this scenario, it is next to impossible to identify the intruder. Suppose this intruder attaches to the WLAN and launches a hacking or denial of service attack against your competitor. The only information you have is an IP address originating from your local network, and if lucky, a MAC address from the DHCP server. By the time you are notified of the attack and gather the information, the attacker could be long gone, never to be traced again.

Positioning is made possible by the fact that certain characteristics measurable in wireless networks vary with respect to the physical location where they measurement is done. This variability is caused by factors such as the distance and the angle between the location of the transmitter and the location of the receiver, and the properties surrounding physical environment, such as the location, shape and material of reflecting and absorbing surfaces like walls.

Unfortunately, measuring these location dependent signals accurately is difficult mainly because of the multi-path problem. The signal measurements are inherently noisy as the radio signals travel between the transmitter and receiver along several alternative paths, and each path is affected by different environmental factors. What is worse, some of the environmental factors are dynamically changing due to presence of humans, variation in air humidity and so on. The situation is especially problematic when there is no line of sight between the transmitter and the receiver as in this case the signal traveling distance is not necessarily the same as the direct signal distance.

The most important factor to decide when tracking is the location-dependent variable that is used. There are a several different methods that can be used to identify and locate unauthorized access points: *time of arrival method, an angle of arrival method and a received signal strength method.*

3.1 Time of Arrival Method (TOA)

It may be possible for the base station to indirectly determine the time that the signal takes from the source to the receiver on the forward or the reverse link. As seen in Figure A.8 this may be done by measuring the time in which the mobile responds to an inquiry or an instruction transmitted to the mobile from the base station. The total time T_m elapsed from the instant the command is transmitted to the instant the mobile response is detected, is composed of the sum of the round trip signal delay $T_1 + T_2$ and any processing and response delay within the mobile unit T_2. If the processing delay for the desired response within the mobile is known with sufficient accuracy, it can be subtracted from total measured time, which would give us the total round trip delay. Half of that quantity would be an estimate of the signal delay in one direction T_p, which would give us the approximate distance of the mobile from the base station. The equations shown below show how to calculate the distance d from the base station to the mobile station.

$$T_m = T_1 + T_2 + T_3 \tag{3.1}$$

$$T_p = \frac{(T_m - T_2)}{2} \tag{3.2}$$

$$d = T_p c \tag{3.3}$$

If timing measurements can be taken at three different base stations then the location of the mobile station can be determined as the unique intersection point of the three circles as seen in Figure A.9.

3.1.1 Advantages

- Can use existing antennas

- Good performance even for low SNR situations

3.1.2 Disadvantages

- Duplex transmission required

- Multipath degrades accuracy

- Response delay in the mobile may vary from different manufacturers

- Susceptible to timing errors in the absence of line of sight conditions

- Requires extremely accurate and synchronized clocks, the shorter the distance between transmitters and receives, the more accurate the clocks must be. A $1\mu s$ error in the clock can lead to a 300 m error in the position.

3.2 Angle of Arrival Method (AOA)

In the angle of arrival (AOA) approach to positioning the location dependent variable is the signal direction, which is the angle at which the signal arrives at a directional antenna at the base stations. A single directional antenna only gives you the bearing, not the distance of the transmitter. Several directional antennas located well apart are required for positioning. For accurate positioning the angle measurements need to be accurate, but achieving high accuracy measurements in wireless networks is difficult because of the multipath problem and other factors.

3.2.1 Advantages

- Works with current mobiles
- Can work with only two base stations
- Passive system

3.2.2 Disadvantages

- Signal measurements are inherently noisy due to the multipath problem and other factors. If one or more of the distance estimates are of poor quality, then the circles drawn around the base stations may not intersect at all
- The location of the base stations needs to be known, otherwise positioning becomes totally impossible
- High accuracy requires expensive directional antenna arrays
- Position estimate degrades as the distance between target and receiver increases
- Accuracy is not very good with existing antennas
- Non-line of sight conditions can result in inaccuracies of 400-700 meters

3.3 Received Signal Strength Method (RSS)

Another technique, is to locate the transmitter using only the relative signal strength at various well placed locations around the area, and if known, the propagation loss, and the power of the transmitter.

If we know the output power of the antenna P_t, the gain of the receiving and transmitting antennas G_r, G_t, the power drop-off, and the received signal strength P_r, then we can determine the location of the transmitter to be one of two possible locations. The equation shown below can be solved for d which will be an estimate of the distance that the mobile station is from the base station. If we don't know the strength of the transmitter or the gain of any of the antennas involved, we just have to introduce one more variable, a signal strength ratio. The problem then turns into a simultaneous equation with three unknowns. To use this method a minimum of three data points is needed.

$$P_r = P_t G_t G_r \left(\frac{\lambda}{4\pi d} \right)^2 \tag{3.4}$$

Assuming two-dimensional geometry, an omni-directional BS antenna, and free-space propagation conditions, signal level contours around a base station are circles. If signal levels from three different BSs are known, the location of the MS can be determined as the unique intersection point of the three circles (view Figure A.11). However, practical propagation conditions especially in urban areas are far from free-space propagation and powerful field prediction tools must be used to compute the signal level as a function of location. Although signal level contours are no longer circles, they can be used for location estimation by finding the location that produces the best fit between predicted and measured values.

3.3.1 Advantages

- Easy to implement
- Low cost positioning method
- Provides relatively good coverage
- Supports existing mobile stations

3.3.2 Disadvantages

- Difficult to achieve high accuracy
- Shadowing has a large effect on accuracy, especially indoors
- A multipath fading environment causes inaccuracies given that the fading characteristics may be different in the direction of the three observation points

3.4 Location Experiments

I intended to perform some location experiments but due to a harddrive crash I was not able to conduct the location experiments. Even though I was not able to get experimental data I will discuss how to perform Angle of Arrival and Received Signal Strength tests.

3.4.1 Angle of Arrival

To use Angle of Arrival you will need a directional antenna such as a parabolic dish. Start by rotating the dish until you get the strongest signal. Record the angle that the dish is pointed. You could use a compass as a means to define your angular coordinate system. You will need to take at least one more measurement at another location. Once you have two or more measurements and know the relative position of where the measurements were taken from you can make a plot with LOBs (lines of bearing). The intersection of the LOBs is the position estimate of the mobile station.

3.4.2 Received Signal Strength

To use the Received Signal Strength method you simply need software which can measure the signal strength of the received signal. Kismet reports the received signal level in dB. You will need to know the power of the transmitter, the power of the received signal, transmitting antenna gain, receiver antenna gain and also an estimate of the propagation conditions (indoor vs. outdoor). By taking 3 or more measurements of the received signal level from different locations you can calculate the distance that the mobile is from each measurement location by using the equation in Section 3.3. You can then plot a circle with a radius of the estimated distance. The three resulting circles should intersect with the mobile station located at the intersection. This method may be the best method to locate wireless APs since it does not require expensive hardware. The RSS method can easily be incorporated into existing wireless APs.

3.5 Location Summary

Each of the location methods has its own advantages and disadvantages. Some issues to consider when choosing a location method are: accuracy, multipath effects, continuous tracking, cost, number of base stations required and if any network modifications are required. A comparison of the location methods can be seen in Table A.2.

If cost is an issue then the Received Signal Strength method would be best because it does not require any special hardware or network modifications except for software. If you require high accuracy a hybrid AOA/RSS or AOA/TOA may be best. A hybrid approach would allow you to use AOA to find the angle of the mobile and the TOA or RSS method would determine how far the mobile is from the base station. In effect you are finding the vector position of the mobile which consists of a distance and an angle.

Chapter 4

Conclusions

This paper examined the problem of rogue access points. A rogue access point can pose a serious security threat to a computer network. If an attacker is able to connect to a unsecured rogue access point they may have access to all of the computers on the inside of the network. The price and ease of use of wireless access points makes it easy for anybody to purchase a wireless access point and simply plug it into the network and have a fully functioning access point, yet the access point is not secure.

Network administrators should routinely scan their networks for rouge access points. Several different methods for detecting rogue access points were discussed. The wired methods for detecting rouge access points are easy because you just have to start a scanning tool and wait for it to return the results. The problems are that it is rather intrusive and may appear to be malicious. Wireless based detection methods are less intrusive and may be faster but they may require the network administrator to walk around the building with a laptop or a wireless handheld device.

Once a rogue access point is detected it is necessary to locate the offending access point so that it can be disabled. As with detection there are also wired and wireless ways to locate a rogue access point. If you can find the MAC address of the access point you can try to determine which network port the AP is connected to. Wireless based methods require 2 or more base stations in order to accurately determine the location. Some methods of position location such as Angle of Arrival require expensive antenna arrays while other methods such as Received Signal Strength require just software modifications. Position location methods suffer from the problems of multipath which can introduce errors into the position estimation. A combination of the above methods may be the best solution. Each wireless access point could be configured to perform RSS location which would give an estimate of the location of the rogue access point. Then the network administrator could use a directional antenna and a laptop to zero in on the rogue access point's position.

Bibliography

[1] S.R. Saunders A.A. Agius. Interference Location Techniques: Final Report. *Radiocommunications Agency Project Reference AY 3639,*, May 2000.

[2] Muhammad Aatique. Evaluation of TDOA Techniques for Position Location In CDMA Systems. *Virginia Polytechnic and State University*, September 1997.

[3] Ekahau Inc. Ekahau Positioning Engine 2.0: 802.11-based Wireless LAN positioning system. *Ekahau Technology Document*, November 2002.

[4] Tero Nordstrom Jaakko Lahteenmaki, Heikki Laiten. Location Methods. *http://location.vtt.fi/source/technologies.html*, April 2001.

[5] Steven A. Parl James M. Zagami. Providing Universal Location Services Using a Wireless E911 Location Network. *Signatron Technology Corporation.*

[6] Eric Junker. Detecting Rogue Access Points. December 2002.

[7] Mark Licht. Wireless 9-1-1 Location: A Positive Move for PSAPs. *911 Magazine*, May 1998.

[8] Interlink Networks. A Practical Approach to Identifying and Tracking Unauthorized 802.11 Cards and Access Points. 2002.

[9] P.K. Chrysanthis P. Prasithsangaree, P. Krishnamurthy. On Indoor Position Location With Wireless LANs. 2002.

[10] Venkata N. Padmanabhan Paramvir Bahl. RADAR: An In-bilding RF-based User Location and Tracking System. *Microsoft Research*, 2000.

[11] Ville Ruutu. Network Location Services. *Nokia Research Center*, April 2003.

[12] Nick Thomas. A Passive Mobile Location Service for UMTS. *University of Edinburgh*, April 1999.

[13] Amy E. Bell Thomas Klein-Ostmann. A Data Fusion Architecture for Enhanced Position Estimation in Wireless Networks. *IEEE Communications Letters*, 5(8), August 2001.

[14] B.D. Woerner T.S. Rappaport, J.H. Reed. Position Location Using Wireless Communications on Highways of the Future. *IEEE Communications*, October 1996.

[15] Joshua Wright. Layer 2 Analysis of WLAN Discovery Applications for Intrusion Detection. November 2002.

[16] Joshua Wright. Detecting Wireless LAN MAC Address Spoofing. January 2003.

Appendix A

Figures

```
# nmap -sT -sU -O 10.0.0.1

Starting nmap V. 2.54BETA29 ( www.insecure.org/nmap )
Interesting ports on (10.0.0.1):
(The 2997 ports scanned but not shown below are in state:  closed)
Port State Service
23/tcp open telnet
68/tcp open bootpc
80/tcp open http
161/udp open snmp

Remote operating system guess:  SonicWall SOHO firewall or Accelerated
Networks VoDSL

Nmap run complete -- 1 IP address (1 host up) scanned in 8 seconds
```

Figure A.1: Nmap scan showing TCP fingerprinting

```
# snmpwalk 10.0.0.1 kkuehl .1.3.6.1.4.1.522.3.1
enterprises.522.3.1.1.0 = "5.4 Hex:  35 2E 34
enterprises.522.3.1.2.0 = "Dec 19 2000, 16:02:33"
enterprises.522.3.1.3.0 = "11.00"
enterprises.522.3.1.4.0 = Guage32:  192
enterprises.522.3.1.5.0 = "en-US"
enterprises.522.3.1.6.0 = Guage32:  3148
enterprises.522.3.1.7.0 = "AP4800-E"
enterprises.522.3.1.8.0 = 1
enterprises.522.3.1.9.0 = 1
enterprises.522.3.1.10.0 = 1
enterprises.522.3.1.11.0 = 1
enterprises.522.3.1.12.0 = 2
enterprises.522.3.1.13.0 = 2
enterprises.522.3.1.14.0 = 2
enterprises.522.3.1.15.0 = Guage32:  64
enterprises.522.3.1.16.0 = 2
enterprises.522.3.1.17.0 = Hex:  00 E0 8F 00 00 00
enterprises.522.3.1.18.0 = "Cisco Systems, Inc."
enterprises.522.3.1.19.0 = "Cisco"
enterprises.522.3.1.20.0 = "http://www.cisco.com"
enterprises.522.3.1.21.0 = 2
enterprises.522.3.1.24.0 = 1
enterprises.522.3.1.25.0 = "1.02" Hex:  31 2E 30 32
```

Figure A.2: SNMP Scanning

	Ease of Use	**Intrusive**	**Continuous**	**Scalability**
TCP Fingerprinting	Easy	Yes	No	Easy
SNMP Scanning	Easy	Yes	No	Easy
Packet Sniffing	Average	No	Yes	Hard
Active Probing	Hard	Yes	Yes	Hard
RF Monitoring	Hard	No	Yes	Hard

Table A.1: Detection Method Comparison

```
[root@optimus]# nmap -sT -sU -O 129.186.8.36

Starting nmap V. 3.00 ( www.insecure.org/nmap/ )
Warning:  OS detection will be MUCH less reliable because we did not find
at least 1 open and 1 closed TCP port
Interesting ports on airport8.design.iastate.edu (129.186.8.36):
(The 3066 ports scanned but not shown below are in state: closed)
Port       State       Service
68/udp     open        dhcpclient
161/udp    open        snmp
192/udp    open        osu-nms
Too many fingerprints match this host for me to give an accurate OS guess

Nmap run completed -- 1 IP address (1 host up) scanned in 11 seconds
```

Figure A.3: TCP Fingerprinting Results

```
[root@optimus]# snmpwalk 129.186.8.36 public system
system.sysDescr.0 = Base Station V3.81 Compatible
system.sysObjectID.0 = OID: .ccitt.zeroDotZero
system.sysUpTime.0 = Timeticks: (49264970) 5 days, 16:50:49.70
system.sysContact.0 = mike miller
system.sysName.0 = Design Airport Station 7 (434)
system.sysLocation.0 = 434 Design
system.sysServices.0 = 79
End of MIB
```

Figure A.4: SNMP Scanning Results

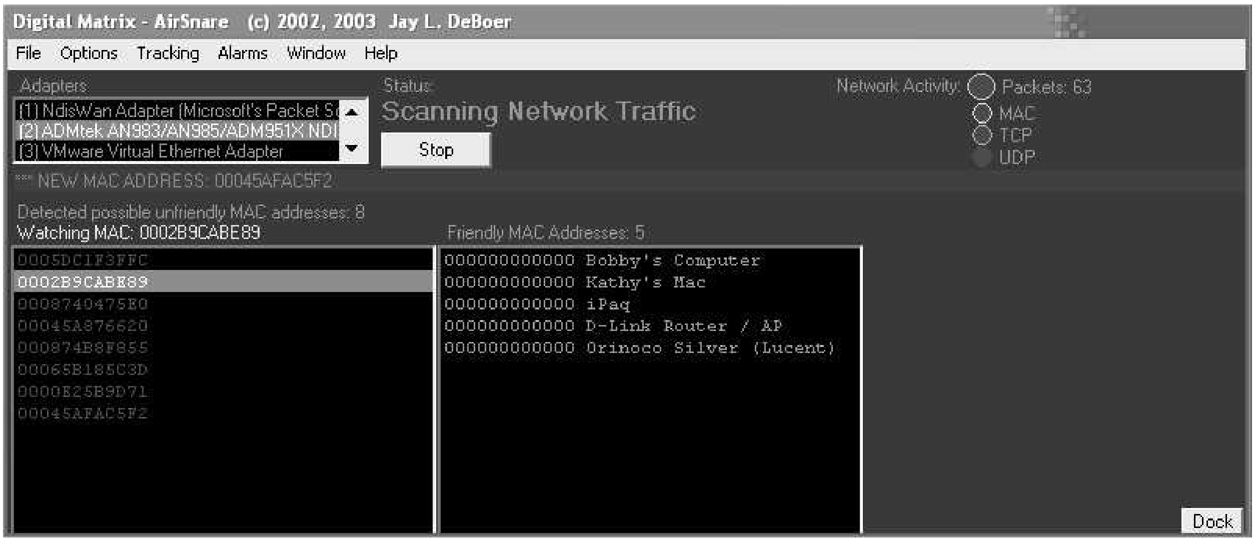

Adapters
```
(1) NdisWan Adapter (Microsoft's Packet Sc
(2) ADMtek AN983/AN985/ADM951X NDI
(3) VMware Virtual Ethernet Adapter
```

Status: Scanning Network Traffic

Network Activity: Packets: 63

```
*** NEW MAC ADDRESS: 00045AFAC5F2
```

Detected possible unfriendly MAC addresses: 8

Watching MAC: 0002B9CABE89

```
0005DC1F3FFC
0002B9CABE89
000874047SE0
00045A876620
000874B8F855
00065B185C3D
0000E25B9D71
00045AFAC5F2
```

Friendly MAC Addresses: 5

```
000000000000 Bobby's Computer
000000000000 Kathy's Mac
000000000000 iPaq
000000000000 D-Link Router / AP
000000000000 Orinoco Silver (Lucent)
```

Figure A.5: Using AirSnare to monitor MAC addresses

	TOA	**AOA**	**RSS**
Accuracy	50-200 m	Up to 50 m	average around 300 m
Multipath Propagation	Large	Large	Small
Continuous Tracking	Yes	Yes	No
Proprietary Hardware	No	Yes	No
Cost	Very High	Very High	Low
Minimum BS Required	3	2	3
Network Modifications	New clocks at each BS	Antenna arrays software	Software

Table A.2: Positioning Technology Comparison

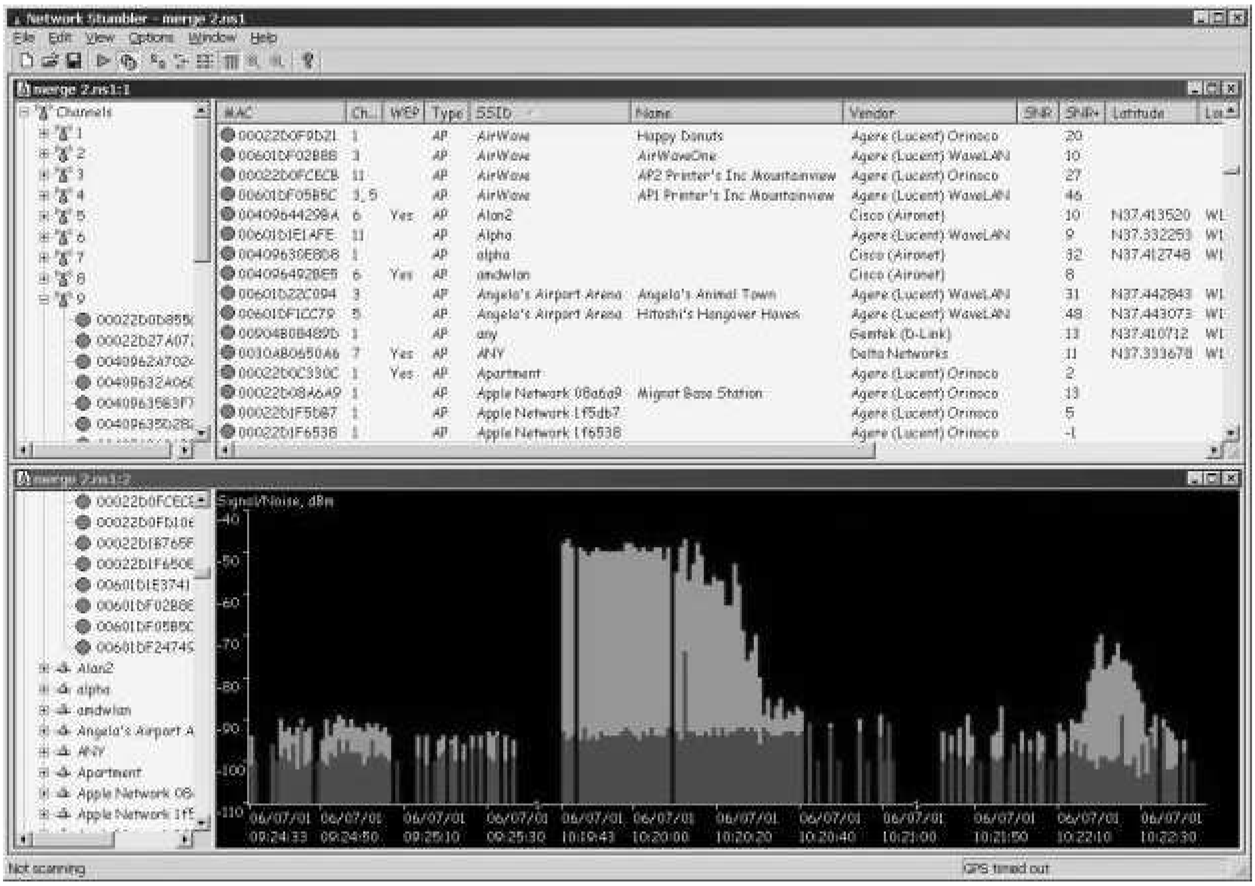

MAC	Ch...	WEP	Type	SSID	Name	Vendor	SNR	SNR+	Latitude	Lo
00022D0F9D21	1		AP	AirWave	Happy Donuts	Agere (Lucent) Orinoco		20		
00601DF02B8B	3		AP	AirWave	AirWaveOne	Agere (Lucent) WaveLAN		10		
00022D0FCECB	11		AP	AirWave	AP2 Printer's Inc Mountainview	Agere (Lucent) Orinoco		27		
00601DF05B5C	1, 5		AP	AirWave	AP1 Printer's Inc Mountainview	Agere (Lucent) WaveLAN		46		
00409644298A	6	Yes	AP	Alan2		Cisco (Aironet)		10	N37.413520	W1
00601D1E1AFE	11		AP	Alpha		Agere (Lucent) WaveLAN		9	N37.332253	W1
00409630E8D8	1		AP	alpha		Cisco (Aironet)		32	N37.412748	W1
004096492BE5	6	Yes	AP	amdwlan		Cisco (Aironet)		8		
00601D22C094	3		AP	Angela's Airport Arena	Angela's Animal Town	Agere (Lucent) WaveLAN		31	N37.442843	W1
00601DF1CC79	5		AP	Angela's Airport Arena	Hitoshi's Hangover Haven	Agere (Lucent) WaveLAN		48	N37.443073	W1
00904B0B489D	1		AP	any		Gemtek (D-Link)		13	N37.410712	W1
00304B0650A6	7	Yes	AP	ANY		Delta Networks		11	N37.333678	W1
00022D0C330C	1	Yes	AP	Apartment		Agere (Lucent) Orinoco		2		
00022D08A6A9	1		AP	Apple Network 08a6a9	Migmot Base Station	Agere (Lucent) Orinoco		13		
00022D1F5D87	1		AP	Apple Network 1f5db7		Agere (Lucent) Orinoco		5		
00022D1F6538	1		AP	Apple Network 1f6538		Agere (Lucent) Orinoco		-1		

Figure A.6: Using Netstumbler to discover wireless APs

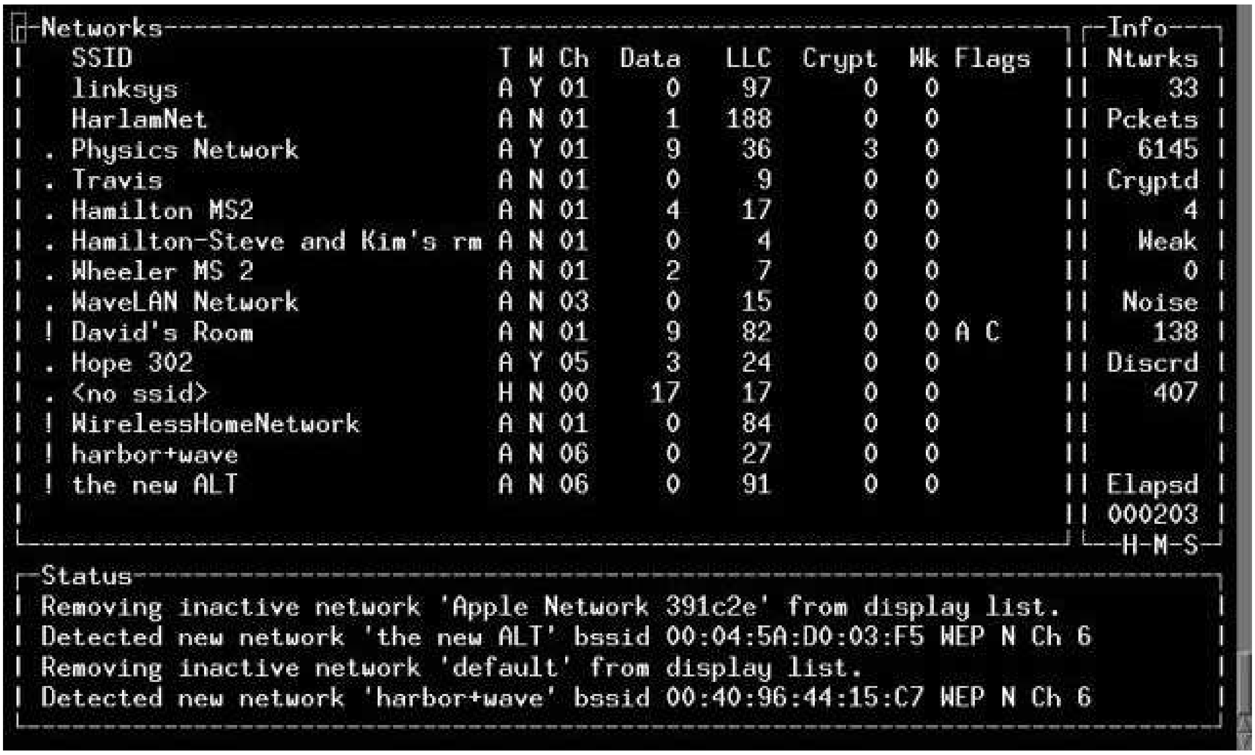

```
┌-Networks-----------------------------------------------------------┐ ┌-Info---┐
|    SSID                          T W Ch  Data   LLC  Crypt  Wk Flags || Ntwrks |
|    linksys                       A Y 01     0    97      0   0       ||     33 |
|    HarlamNet                     A N 01     1   188      0   0       || Pckets |
| . Physics Network                A Y 01     9    36      3   0       ||   6145 |
| . Travis                         A N 01     0     9      0   0       || Cryptd |
| . Hamilton MS2                   A N 01     4    17      0   0       ||      4 |
| . Hamilton-Steve and Kim's rm    A N 01     0     4      0   0       ||   Weak |
| . Wheeler MS 2                   A N 01     2     7      0   0       ||      0 |
| . WaveLAN Network                A N 03     0    15      0   0       ||  Noise |
| ! David's Room                   A N 01     9    82      0   0 A C   ||    138 |
| . Hope 302                       A Y 05     3    24      0   0       || Discrd |
| . <no ssid>                      H N 00    17    17      0   0       ||    407 |
| ! WirelessHomeNetwork            A N 01     0    84      0   0       ||        |
| ! harbor+wave                    A N 06     0    27      0   0       ||        |
| ! the new ALT                    A N 06     0    91      0   0       || Elapsd |
|                                                                      || 000203 |
└----------------------------------------------------------------------┘ └-H-M-S-┘
┌-Status----------------------------------------------------------------------------┐
| Removing inactive network 'Apple Network 391c2e' from display list.               |
| Detected new network 'the new ALT' bssid 00:04:5A:D0:03:F5 WEP N Ch 6             |
| Removing inactive network 'default' from display list.                            |
| Detected new network 'harbor+wave' bssid 00:40:96:44:15:C7 WEP N Ch 6             |
└-----------------------------------------------------------------------------------┘
```

Figure A.7: Using Kismet to discover wireless APs

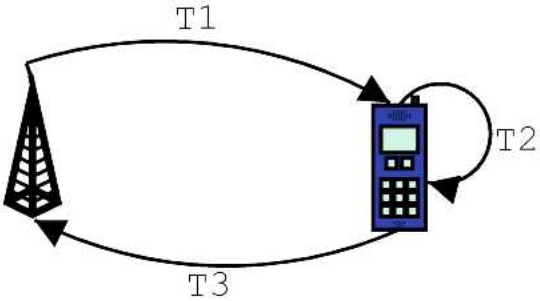

Figure A.8: Time measurements

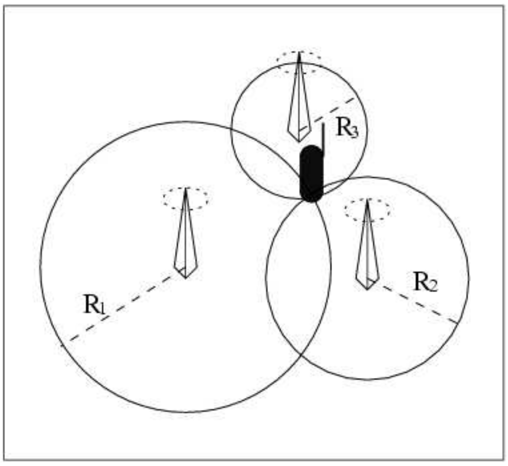

Figure A.9: Positioning using time of arrival

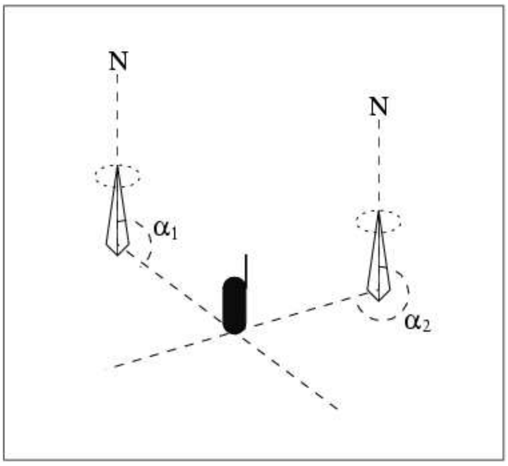

Figure A.10: Positioning using angle of arrival

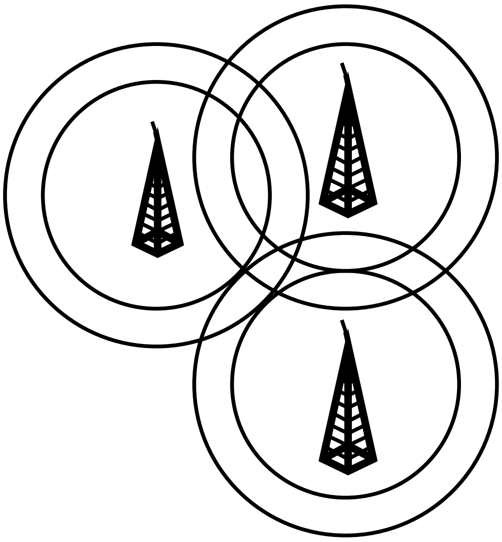

Figure A.11: Tracking by signal strength